arXiv:1704.03503v1 [cs.CV] 11 Apr 2017

UC Merced Submission to the ActivityNet Challenge 2016

Yi Zhu Shawn Newsam
University of California, Merced, USA
{yzhu25,snewsam}@ucmerced.edu

Zaikun Xu
University of Lugano, Switzerland
xuz@usi.ch

Abstract

This notebook paper describes our system for the untrimmed classification task in the ActivityNet challenge 2016. We investigate multiple state-of-the-art approaches for action recognition in long, untrimmed videos. We exploit hand-crafted motion boundary histogram features as well feature activations from deep networks such as VGG16, GoogLeNet, and C3D. These features are separately fed to linear, one-versus-rest support vector machine classifiers to produce confidence scores for each action class. These predictions are then fused along with the softmax scores of the recent ultra-deep ResNet-101 using weighted averaging.

1. Introduction

Human action recognition in video is a fundamental problem in computer vision due to its increasing importance for a range of applications such as video recommendation and search, video highlighting, video surveillance, human-robot interaction, human skill evaluation, etc.

The ActivityNet challenge [4] is a large scale benchmark designed to stimulate research on human activity understanding in user generated videos. This challenge consists of two tasks on 200 activity categories: (a) untrimmed classification and (b) detection. We focus on the former which involves predicting the activities present in a long video. Accounting for YouTube blocks and deleted videos, we downloaded 9942 training, 4874 validation, and 5001 test videos.

2. Recognition Framework

In this section, we present our multi-stream action recognition framework based on: (i) Fisher vector encoded MBH features, (ii) C3D fc7 features, (iii) GoogLeNet pool5 features, (iv) VGG16 pool5 features, and (v) ResNet-101 softmax scores. The first two modules are clip-based while the last three are frame-based. An overview of the framework can be found in Fig. 1.

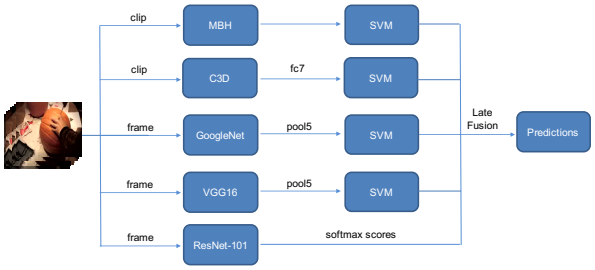

Figure 1. **Multi-stream framework**. We combine five modules using late fusion to obtain the final prediction scores. The MBH module is hand-crafted, while the rest are based on deep networks. For ResNet-101, we directly use the softmax scores since this performs better than using the extracted features.

2.1. MBH Features

Improved dense trajectories (IDT) [15] are state-of-the-art hand-crafted features for modeling temporal information in videos, and the motion boundary histogram (MBH) features are the best performing component of the IDT features. We use the provided[1] Fisher vector encoded MBH features [13, 9], whose dimension is 65536 for each video, to train a linear, one-versus-rest support vector machine (SVM) classifier. We fix the SVM hyper-parameter C to 100 [2].

2.2. C3D

In [14], the authors show that 2D ConvNets "forget" the temporal information in the input signal after each convolution. They therefore propose 3D ConvNets, which analyze sets of contiguous video frames organized as clips, and show its effectiveness at learning spatio-temporal features in video volume data analysis problems.

We therefore adopt fc7 features[2] extracted from a pre-trained C3D model as an additional signal. The network is not fine-tuned on the ActivityNet challenge dataset. The inputs to the C3D model are 16 frame clips with 50% overlap and the outputs are 4096 dimension feature activations.

[1] The MBH features are provided by the organizers.
[2] The C3D extracted fc7 features are provided by the organizers.

These features are reduced to 500 dimensions using PCA. Average pooling is used to combine the clip-level features for a single video. Finally, a linear, one-versus-rest SVM is trained with C set to 1.

2.3. GoogLeNet

We also use features[3] extracted from the pool5 layer of a Google inception net (GoogLeNet) [8]. This network is an enhanced version of [12] which utilizes a reorganized hierarchy of the complete ImageNet dataset [1]. The features are frame-based and have dimension 1024. They are mean-pooled across all frames in a video followed by L1-normalization to obtain a video-level representation. Again, a linear, one-versus-rest SVM is trained with C set to 1.

2.4. VGG16

VGG16 [11] is a popular deep architecture that demonstrated good performance on action recognition in [16] using a two-stream [10] pipeline. We only employ the spatial stream. We use the pre-trained VGG16 model for initialization and fine-tune it on the challenge dataset. During fine-tuning, we perform 60K iterations with learning rate 10^{-4}, 30K iterations with 10^{-5}, and 30K iterations with 10^{-6}. Momentum and weight decay are set to 0.9 and 5×10^{-4}.

We adopt the latent concept descriptor (LCD) encoding method in [17] to encode the pool5 layer of our fine-tuned VGG16 model, followed by VLAD encoding [6]. We reduce the dimensions of the pool5 features from 512 to 256 using PCA. The number of centers in VLAD encoding is set to 256 and we use VLAD-k with k set to 5. The encoded features are then power- and intra-normalized. The resulting per-frame features have dimension 65536 which are mean-pooled to obtain a video-level representation. A linear, one-versus-rest SVM is trained with C set to 1.

2.5. ResNet-101

Residual learning [3] has recently become an effective method to construct ultra-deep networks for object recognition and detection. We extend it here to action recognition. We adopt the pre-trained 101-layer model for initialization and fine-tune it on the ActivityNet video data. The learning rate is 10^{-4} for the first two epochs, 10^{-5} for the following two epochs, and 10^{-6} for the last epoch. Momentum and weight decay are set to 0.9 and 10^{-4}.

We also investigated using features extracted from last convolutional module, whose dimension is 2048, to train an SVM, similar to our other modules. This, however, performs 3.3% worse on the validation set than using the softmax scores.

[3]The GoogLeNet extracted pool5 features are provided by the organizers.

Model	Top-1 Accuracy
(i) MBH	57.32%
(ii) C3D fc7	60.04%
(iii) GoogLeNet pool5	67.13%
(iv) VGG16* pool5	63.19%
(v) ResNet-101*	71.81%
(i) + (ii)	62.78%
(i) + (iii)	69.40%
(i) + (iv)	68.79%
(ii) + (iii)	68.11%
(ii) + (iv)	64.35%
(iii) + (iv)	68.56%
(ii) + (iii) + (iv)	69.09%
(i) + (v)	73.05%
(ii) + (iii) + (iv) + (v)	73.56%
(i) + (iii) + (iv) + (v)	74.68%
(i) + (ii) + (iii) + (iv) + (v)	75.14%

Table 1. Action recognition results on the validation set of the ActivityNet challenge 2016. All performances are reported using top-1 accuracy. Top: Single module performances. Bottom: Fused module performances. * indicates the network is fine-tuned on the challenge dataset.

3. Experiment Results

Given a test video, we uniformly sample 25 frames to extract the frame-level feature activations and perform mean-pooling to obtain the final video representation.

Late fusion iteratively combines pairs of prediction scores. First, the outputs of two modules are combined in a weighted fashion where the scores of the more accurate module are weighted twice that of the less accurate one. Additional scores are then combined with this in a similar fashion. After late fusion, we adopt a Multi-class Iterative Re-ranking (MIR) method [7] to re-rank the predictions of classifiers based on the difficulty scale of the videos. Table 1 shows our experimental results on the validation set of the ActivityNet challenge 2016.

We can see from the top part of Table 1 that the residual network achieves the best performance among all modules. It is 14.5% better than the state-of-the-art hand-crafted MBH features and outperforms the other deep networks.

The bottom part of Table 1 shows the performances of various module combinations. We observe that combinations that only include deep networks are generally not as effective as combinations that include the MBH features. Although the MBH features perform the worst alone, they are orthogonal to the deep learning based approaches. This may be attributed to MBH being effective at capturing low-level motion features while the deep networks model high-

Submission	mAP	Top-1 Accuracy	Top-3 Accuracy
Run 1	68.00%	66.16%	83.36%
Run 2	75.98%	72.48%	87.54%
Run 3	79.41%	76.17%	90.19%
Run 4	81.64%	77.74%	90.93%
Run 5	83.1%	78.44%	91.07%

Table 2. Action recognition results on the test set of the ActivityNet challenge 2016.

level information related to static appearance. The MBH features and the deep networks are thus quite complementary. When fusing all modules, our system achieves a validation accuracy of 75.14%.

We also investigate incorporating action proposals generated by [5] during prediction. Instead of uniformly sampling 25 frames across the video, we sample 25 frames from the action proposals. The intuition is that these action proposals have a higher probability of containing action frames, so that the average pooling of these frames should lead to higher recognition accuracy. However, this turns out to perform worse than uniform sampling.

4. Submission Details

We use both the training and validation data as the training set for our submissions. Note, though, that the implementation details and parameter settings remain the same as when we use only the training data to train. We do not use the test data for training or parameter tuning.

We submit five runs to the evaluation server, and the performance for each run is shown in Table 2. Our runs are as follows:

- Run 1: VGG16
- Run 2: VGG16 + MBH
- Run 3: VGG16 + MBH + ResNet-101
- Run 4: VGG16 + MBH + ResNet-101 + GoogLeNet
- Run 5: VGG16 + MBH + ResNet-101 + GoogLeNet + C3D

5. Conclusion

We show that the ultra-deep architecture of ResNet is indeed helpful in learning discriminative features for complex tasks, such as human activity understanding. In addition, although hand-crafted MBH features achieve the lowest accuracy alone, they are complementary to approaches based on deep networks. Finally, the combination of all modules using late fusion gives the best performance.

Acknowledgement

We gratefully acknowledge the support of NVIDIA Corporation through the donation of the Titan X GPU used in this work. This work was funded in part by a National Science Foundation CAREER grant, #IIS-1150115.

References

[1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 2

[2] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A Library for Large Linear Classification. *Journal of Machine Learning Research*, 2008. 1

[3] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 2

[4] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding. In *CVPR*, 2015. 1

[5] F. C. Heilbron, J. C. Niebles, and B. Ghanem. Fast Temporal Activity Proposals for Efficient Detection of Human Actions in Untrimmed Videos. In *CVPR*, 2016. 3

[6] H. Jegou, F. Perronnin, M. Douze, J. Sanchez, P. Perez, and C. Schmid. Aggregating Local Image Descriptors into Compact Codes. *TPAMI*, 2012. 2

[7] Z. Lan, S.-I. Yu, and A. G. Hauptmann. Improving Human Activity Recognition Through Ranking and Re-ranking. *arXiv preprint arXiv:1512.03740*, 2015. 2

[8] P. Mettes, D. Koelma, and C. G. M. Snoek. The ImageNet Shuffle: Reorganized Pre-training for Video Event Detection. In *ICMR*, New York, USA, 2016. 2

[9] D. Oneata, J. Verbeek, and C. Schmid. Action and Event Recognition with Fisher Vectors on a Compact Feature Set. In *ICCV*, 2013. 1

[10] K. Simonyan and A. Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. *NIPS*, 2014. 2

[11] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2015. 2

[12] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going Deeper With Convolutions. In *CVPR*, 2015. 2

[13] J. Snchez, F. Perronnin, T. Mensink, and J. Verbeek. Image Classification with the Fisher Vector: Theory and Practice. *IJCV*, 2013. 1

[14] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning Spatiotemporal Features with 3D Convolutional Networks. In *ICCV*, 2015. 1

[15] H. Wang and C. Schmid. Action Recognition with Improved Trajectories. In *ICCV*, 2013. 1

[16] L. Wang, Y. Xiong, Z. Wang, and Y. Qiao. Towards Good Practices for Very Deep Two-Stream ConvNets. *arXiv preprint arXiv:1507.02159*, 2015. 2

[17] Z. Xu, Y. Yang, and A. G. Hauptmann. A Discriminative CNN Video Representation for Event Detection. In *CVPR*, 2015. 2